09451 Abstracts Collection

Geometric Networks, Metric Space Embeddings and Spatial Data Mining

— Dagstuhl Seminar —

Gautam Das,[1], Joachim Gudmundsson[2], Rolf Klein[3], Christian Knauer[4] and Michiel Smid[5]

[1] University of Texas at Arlilngton, US
gdas@cse.uta.edu
[2] National ICT Australia - Sydney, AU
joachim.gudmundsson@nicta.com.au
[3] Universität Bonn, D
rolf.klein@uni-bonn.de
[4] FU Berlin, D
knauer@inf.fu-berlin.de
[5] Carleton Univ. - Ottawa, CDN
michiel@scs.carleton.ca

Abstract. From November 1 to 6, 2009, the Dagstuhl Seminar 09451 “Geometric Networks, Metric Space Embeddings and Spatial Data Mining” was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Geometric networks, metric space embeddings, spatial data mining, spanners, dilation, distortion

Dagstuhl Seminar Proceedings 09451
Geometric Networks, Metric Space Embeddings and Spatial Data Mining
http://drops.dagstuhl.de/opus/volltexte/2010/2438

1 Summary of the seminar

This seminar has brought together scientists from three different communities (geometric networks, metric space embeddings, and spatial data mining) who have numerous interests in common, all of which are related to distance problems. The seminar was a continuation of the Dagstuhl seminar 06481 (Geometric Networks and Metric Space Embeddings) which was held in 2006. The main purpose of the current seminar was to continue, and intensify, the cooperation between the geometric network and the metric space communities. This time, we invited people from spatial data mining to provide the extra application stimulus.

A *geometric network* is a graph mapped into the Euclidean plane or a Euclidean space of low dimension. It is said to be a *spanner* if the network distance between any pair of nodes is bounded by a constant times their Euclidean distance. Geometric networks are the backbone of any model for the flow of goods, traffic or information. They also play an important role in telecommunication, VLSI design, motion planning (robotics), pattern matching, data compression, bio-informatics (gene analysis), and sensor networks. One is interested in spanners with other useful properties such as a linear number of edges, small total edge length, small node degree, few crossings, or small link diameter. Apart from these applications, geometric spanners have had great impact on the construction of approximation algorithms, e.g., for the traveling salesperson problem.

The similarity between individual objects of a given finite domain becomes apparent if the objects can be represented by points in the plane, or in 3-space, in such a way that the Euclidean distance between two points corresponds to the similarity of the objects they are associated with. The question when such representations exist has led to the theory of embedding finite metric spaces into normed vector spaces. It is of particular interest for storage, visualization, and retrieval of high-dimensional data, e.g., in information retrieval.

Both problems (metric space embedding and spanner construction) have received a lot of attention during the last 10-15 years, within their respective communities. Indeed, the first monograph on spanners (co-authored by the 5th organizer) has been published meanwhile, and metric space embeddings have been addressed in several books and book chapters. In both cases, we are applying two different metrics to the point pairs of the same set, and we are looking for the maximum (or minimum) ratio of all values. In metric space embeddings we compare the measure of similarity of the objects against the (Euclidean) distance of their associated points, and the maximum ratio is called the distortion of the embedding. In a spanning geometric network, we compare the shortest path distance in the network against the Euclidean distance between two points; here, the extremal ratio is called dilation or stretch factor. In both areas many questions are open. For example, it is not known how to construct the triangulation of minimum dilation over a given point set, with or without extra Steiner points allowed.

Data mining can be seen as the science of extracting useful information from large data sets or databases. It is the principle of sorting through large amounts of data and finding relevant information. It is usually used by business intelligence organizations and financial analysts, but it is increasingly used in the sciences to extract information from the enormous data sets generated by modern experimental and observational methods. In many applications of data mining, the high dimensionality of the data restricts the choice of data processing methods. Such application areas include the analysis of market basket data, text documents, image data and so on; in these cases the dimensionality is large due to either a wealth of alternative products, a large vocabulary, or the use of large image windows, respectively. A common tool used to develop efficient algorithms is to reduce the number of dimensions. A statistically optimal way of dimensionality reduction is to project the data onto a lower-dimensional orthogonal subspace that captures as much of the variation of the data as possible. The best (in mean-square sense) and most widely used way to do this is principal component analysis (PCA); unfortunately, it is quite expensive to compute for high-dimensional data sets. A computationally simple method of dimensionality reduction that does not introduce a signifficant distortion in the data set is random projection. Here, the original high-dimensional data is projected onto a lower-dimensional subspace using a random matrix whose columns have unit lengths. Random projection has been found to be a computationally efficient, yet sufficiently accurate method for dimensionality reduction of high-dimensional data sets.

The seminar's aim was at crossfertilization between the three communities. The main results of the seminar can be summarized as follows.

- During the seminar, it became clear that methods developed in the theory of finite metric spaces help in analyzing geometric networks. Conversely, the algorithmic techniques developed in geometric network design are of interest to people working on embedding problems.
- There was a fruitful exchange of ideas which stimulated interesting discussions and future cooperations.
- The seminar will advance a comparative theory of distance measures.
- The knowledge gained during the seminar will help in reducing the complexity of high-dimensional data, as is important in data mining and related areas.

2 List of talks and abstracts

Geometric Spanners for Weighted Point Sets

Mohammad Ali Abam (Aarhus University - Aarhus, DK)

Let $(S, \mathbf{d})$ be a finite metric space, where each element $p \in S$ has a non-negative weight $\mathrm{w}(p)$. We study spanners for the set S with respect to weighted distance function $\mathbf{d}_w$, where $\mathbf{d}_w(p,q)$ is $\mathrm{w}(p) + \mathbf{d}(p,q) + \mathrm{w}(q)$ if $p \neq q$ and 0 otherwise.

We present a general method for turning spanners with respect to the $\mathbf{d}$-metric into spanners with respect to the $\mathbf{d}_w$-metric. For any given $\epsilon > 0$, we can apply our method to obtain $(5+\epsilon)$-spanners with a linear number of edges for three cases: points in Euclidean space $\mathbb{R}^d$, points in spaces of bounded doubling dimension, and points on the boundary of a convex body in $\mathbb{R}^d$ where $\mathbf{d}$ is the geodesic distance function.

We also describe an alternative method that leads to $(2+\epsilon)$-spanners for points in $\mathbb{R}^d$ and for points on the boundary of a convex body in $\mathbb{R}^d$. The number of edges in these spanners is $O(n \log n)$. This bound on the stretch factor is nearly optimal: in any finite metric space and for any $\epsilon > 0$, it is possible to assign weights to the elements such that any non-complete graph has stretch factor larger than $2-\epsilon$.

Joint work with: de Berg, Mark; Farshi, Mohammad; Gudmundsson, Joachim; Smid, Michiel

Kinetic Spanner in $\mathbb{R}^d$

Mohammad Ali Abam (Aarhus University - Aarhus, DK)

We present a new $(1+\epsilon)$-spanner for sets of n points in $\mathbb{R}^d$. Our spanner has size $O(n/\epsilon^{d-1})$ and maximum degree $O(\log^d n)$. The main advantage of our spanner is that it can be maintained efficiently as the points move: Assuming the trajectories of the points can be described by bounded-degree polynomials, the number of topological changes to the spanner is $O(n^2/\epsilon^{d-1})$, and using a supporting data structure of size $O(n \log^d n)$ we can handle events in time $O(\log^{d+1} n)$. Moreover, the spanner can be updated in time $O(\log n)$ if the flight plan of a point changes. This is the first kinetic spanner for points in $\mathbb{R}^d$ whose performance does not depend on the spread of the point set.

Joint work with: de Berg, Mark

On Plane Geometric Spanners

Prosenjit Bose (Carleton University - Ottawa, CA)

We review results and present open problems on different variants of the problem of constructing plane geometric spanners.

Keywords: Spanner, Geometric Graph

Joint work with: Smid, Michiel

The Fibonacci Dimension of a Graph

Sergio Cabello (University of Ljubljana, SI)

The Fibonacci dimension $fdim(G)$ of a graph G is introduced as the smallest integer f such that G admits an isometric embedding into Γ_f, the f-dimensional Fibonacci cube.

We give bounds on the Fibonacci dimension of a graph in terms of the isometric and lattice dimension, provide a combinatorial characterization of the Fibonacci dimension using properties of an associated graph, and establish the Fibonacci dimension for certain families of graphs.

From the algorithmic point of view we prove that it is NP-complete to decide if $fdim(G)$ equals to the isometric dimension of G, and that it is also NP-hard to approximate $fdim(G)$ within $(741/740) - \epsilon$. We also give a $(3/2)$-approximation algorithm for $fdim(G)$ in the general case and a $(1+\epsilon)$-approximation algorithm for simplex graphs.

Joint work with: Eppstein, David; Klavzar, Sandi

Full Paper:
http://arxiv.org/abs/0903.2507

On the Power of the Semi-Separated Pair Decomposition

Paz Carmi (Ben Gurion University - Beer Sheva, IL)

A Semi-Separated Pair Decomposition (SSPD), with parameter $s > 1$, of a set $S \subset \mathbb{R}^d$ is a set $\{(A_i, B_i)\}$ of pairs of subsets of S such that for each i, there are balls D_{A_i} and D_{B_i} containing A_i and B_i respectively such that $d(D_{A_i}, D_{B_i}) \geq s \cdot \min(\text{radius}(D_{A_i}), \text{radius}(D_{B_i}))$, and for any two points $p, q \in S$ there is a unique index i such that $p \in A_i$ and $q \in B_i$ or vice-versa.

In this talk, we use the SSPD to obtain the following results:

First, we consider the construction of geometric t-spanners in the context of imprecise points and we prove that any set $S \subset \mathbb{R}^d$ of n imprecise points, modeled as pairwise disjoint balls, admits a t-spanner with $O(n \log n/(t-1)^d)$ edges which can be computed in $O(n \log n/(t-1)^d)$ time. If all balls have the same radius, the number of edges reduces to $O(n/(t-1)^d)$.

Secondly, for a set of n points in the plane, we design a query data structure for half-plane closest-pair queries that can be built in $O(n^2 \log^2 n)$ time using $O(n \log n)$ space and answers a query in $O(n^{1/2+\epsilon})$ time, for any $\epsilon > 0$. By reducing the preprocessing time to $O(n^{1+\epsilon})$ and using $O(n \log^2 n)$ space, the query can be answered in $O(n^{3/4+\epsilon})$ time. Moreover, we improve the preprocessing time of an existing axis-parallel rectangle closest-pair query data structure from quadratic to near-linear.

Finally, we revisit some previously studied problems, namely spanners for complete k-partite graphs and low-diameter spanners, and show how to use the SSPD to obtain simple algorithms for these problems.

Joint work with: Abam, Mohammad Ali; Farshi, Mohammad; Smid, Michiel

Outlier Detection

Sanjay Chawla (The University of Sydney - Sydney, AU)

An outlier is an unexpected event or entity. For example, many experts view the Great Financial Crash (GFC) of 2008 as an outlier event which has triggered a reappraisal of mainstream economic and financial models which define the "expected." The objective of Outlier Detection in Data Mining is in similar vein — outliers often embody new information, which is often hard to explain in the context of existing knowledge and results in a re-evaluation of what is known. In this talk I will begin by presenting a brief overview of modern outlier detection techniques. In particular I will focus on global and local outlier detection methods and then propose a new approach, based on spectral embedding, which unifies both these strands.

Isometric Embeddings of Cube-Free Median Graphs into Cartesian Products of Trees

Victor Chepoi (Université de la Méditerrannée - Marseille, FR)

Median graphs and related median structures (median algebras and median complexes) have many nice properties and admit numerous characterizations. These structures have been investigated in several contexts (combinatorics, graph theory, universal algebra, geometric group theory, geometry of spaces of nonpositive curvature, operation research) by quite a number of authors for more than half a century. Median graphs represent one of the most important graph classes in metric graph theory.

It is well-known that median graphs isometrically embeds into hypercubes and therefore into Cartesian products of trees. On the other hand, computing the least number of tree factors is NP-hard (even in the case of cube-free median graphs) by a reduction from GRAPH COLORING.

In this talk, after an overview of characterizations and properties of median graphs and complexes, we will give a structural characterization of all graphs which can be isometrically into the Cartesian product of two trees as the cube-free median graphs not containing odd wheels, i.e., cube-free median graphs having bipartite links, and which we call partial double trees. The rectangular complexes derived from partial double dendrons and endowed with intrinsic l_1-metric are shown to be embeddable into the cartesian product of two dendrons.

To establish the first result, we characterize median graphs which can be embedded into the Cartesian product of n trees as the median graphs whose edges can be colored with n colors so that the incident edges of each 4-cycle get different colors and opposite edges get the same color.

A particular class of cube-free median graph is constituted by the square-graphs: these are the plane graphs in which all inner faces are quadrilaterals (i.e., 4-cycles) and the degrees of all inner vertices (i.e., the vertices not incident with the outer face) have degrees larger than three. We establish that any squaregraph can be isometrically embedded into the cartesian product of at most five trees.

For this, we show that the planar dual of a finite squaregraph is determined by a triangle-free chord diagram of the unit disk, which could alternatively be viewed as a triangle-free line arrangement in the hyperbolic plane. This representation carries over to infinite plane graphs with finite vertex degrees in which the balls are finite squaregraphs. Algebraically, finite squaregraphs are median graphs for which the duals are finite circular split systems. Hence squaregraphs are at the crosspoint of two dualities, an algebraic and a geometric one, and thus lend themselves to several combinatorial interpretations and structural characterizations. With these and the 5-colorability theorem for circle graphs at hand, we prove that every squaregraph can be isometrically embedded into the Cartesian product of five trees. This embedding result can also be extended to the infinite case without reference to an embedding in the plane and without any cardinality restriction when formulated for median graphs free of cubes and further finite obstructions.

Keywords: Isometric embedding, Cartesian product of trees, median graph

Joint work with: Hans-Juergen Bandelt; David Eppstein

Top-k Queries Using Windowing

Gautam Das (University of Texas - Arlington, US)

The problem of obtaining efficient answers to top-k queries has attracted a lot of research attention. Several algorithms and numerous variants of the top-k retrieval problem have been introduced in recent years. The general form of this problem requests the k highest ranked values from a relation, using monotone combining functions on (a subset of) its attributes. In this paper we explore space performance tradeoffs related to this problem. In particular we study the problem of answering top-k queries using views. A view in this context is a materialized version of a previously posed query, requesting a number of highest ranked values according to some monotone combining function defined on a subset of the attributes of a relation. Several problems of interest arise in the presence of such views. We start by presenting a new algorithm capable of combining the information from a number of views to answer ad hoc top-k queries. We then address the problem of identifying the most promising (in terms of performance) views to use for query answering in the presence of a collection of views. We formalize both problems and present efficient algorithms for their solution. We also discuss several extensions of the basic problems in this setting. We present the results of a thorough experimental study that deploys our techniques on real and synthetic data sets. Our results indicate that the techniques proposed herein comprise a robust solution to the problem of top-k query answering using views, gracefully exploring the space versus performance tradeoffs in the context of top-k query answering.

Metric Spanning Trees

Michael Elkin (Ben Gurion University - Beer Sheva, IL)

Given a metric M or a graph G one often wants to build a spanning tree T for M or G that satisfies some useful properties. In particular, we may want T to have small weight, or low average stretch, or low hop-diameter. In this talk, we will overview the emerging theory of spanning trees that satisfy one or more of these properties.

Specifically, we plan to start with describing the seminal works of Alon, Karp, Peleg, and West, and of Bartal, and then turn to the more recent work by Emek, Spielman, Teng and the speaker. We will also discuss the minimum total distance tree problem (Wong, 80). Finally, we will overview the constructions of shallow-light and low-light trees.

Keywords: Metric embeddings, spanning trees

Spectral Partitioning

Joachim Giesen (Universität Jena, DE)

A partitioning of a set of n items is a grouping of these items into k disjoint, equally sized classes. Any partition can be modeled as a graph. The items become the vertices of the graph and two vertices are connected by an edge if and only if the associated items belong to the same class. In a planted partition model, a graph that models a partition is given, which is obscured by random noise, i.e., edges within a class can get removed and edges between classes can get inserted. The task is to reconstruct the planted partition from this graph. We design a spectral partitioning algorithm

Keywords: Stewart Theorem, Furedi-Komlos Theorem

Joint work with: Dieter Mitsche

See also: LNCS Volume 3787/2005

Shortest Paths using Currents

Joachim Gudmundsson (NICTA - Sydney, AU)

Ocean currents have been steering the course of sailing vessels for centuries and they are an integral part of ship navigation.

This is an attempt to model navigation of vessels among ocean currents while minimizing the cost of the trip. In this talk, an ocean current is represented by directed straight line segments with an associated cost function, which reflects the strength of the current. The induced metric is called a transportation metric.

We give an exact algorithm for the basic version of the problem: given a set of currents, a source point s and a destination point t, find the cheapest path between s and t among the currents. We also show how the set of currents can be preprocessed in time $O(n^4)$ into a data structure of size $O(n^2)$ such that $(1+\epsilon)$-approximate cheapest path queries between any two points in the plane can be answered in time $O(\log n + k)$, where k is the complexity of the reported path.

A Traveller's Problem

Rolf Klein (Universität Bonn, DE)

A traveller is planning a tour from some start position s to a goal position g in d-dimensional space. Transportation is provided by n carriers. Each carrier is a convex object that results from intersecting finitely many closed linear subspaces; it moves at constant speed along a line. Different carriers may be assigned different velocity vectors. While using carrier C, the traveller can walk at eigenspeed $v \geq 0$ in any direction, like a passenger on board a vessel. Whenever his current position on C is simultaneously contained in some other carrier C', the traveller can change from C to C', and continue his tour by C'.

Given initial positions of the carriers and of s and g, is the traveller able to reach g starting from s? If so, what minimum travel time can be achieved? We provide the following answers.

For a situation similar to the "Frogger" game, where the traveller has to cross a river on which n consecutive rectangular barges move at m different speeds, we provide an $O(n \log m)$ solution.

In dimension 8 and higher, the traveller's problem is undecidable, even for eigenspeed zero.

An interesting case is in dimension 2. We prove that the problem is NP-hard, even if all carriers are vertical line segments. It turns out that an s-to-g path of finite duration may require an infinite number of carrier changes. Despite this difficulty, we can show that the two-dimensional problem is decidable. In addition, we provide a pseudo-polynomial approximation algorithm.

Keywords: Affine mappings, computational geometry, continuous Dijkstra, frogger, motion planning, NP-hardness, partition, pseudo-polynomial approximation, undecidability

Joint work with: Berger, Florian

On the Optimality of Spiral Search

Elmar Langetepe (Universität Bonn, DE)

Searching for a point in the plane is a well-known search game problem introduced in the early eighties. The best known search strategy is given by a spiral and achieves a *competitive ratio* of 17.289... It was shown by Gal in 1980 that this strategy is the best strategy among all *monotone* and *periodic* strategies.

Since then it was unknown whether the given strategy is optimal in general. This talk settles this old open fundamental search problem and shows that spiral search is indeed optimal. The given problem can be considered as the continuous version of the well-known m-ray search problem and also appears in several non-geometric applications and modifications. Therefore the optimality of spiral search is an important question considered by many researchers in the last decades. We answer the logarithmic spiral conjecture for the given problem. The lower bound construction might be helpful for similar settings, it also simplifies existing proofs on classical m-ray search.

Keywords: Search games, computational geometry, motion planning, spiral search, m-ray search, competitive analysis, lower bound

See also: Proc. 21st Annu. ACM-SIAM Symp. Disc. Algor., 2010.

Optimal Randomized Algorithm for the Density Selection Problem

Tien-Ching Lin (Academia Sinica - Taipei, TW)

In the talk we consider a generalized version of three well-known problems: SELECTION PROBLEM in computer science, SLOPE SELECTION PROBLEM in computational geometry and MAXIMUM-DENSITY SEGMENT PROBLEM in bioinformatics. Given a sequence $A = (a_1, w_1), (a_2, w_2), \ldots, (a_n, w_n)$ of n ordered pairs (a_i, w_i) of real numbers a_i and $w_i > 0$ for each $1 \leq i \leq n$, two nonnegative real numbers ℓ, u with $\ell \leq u$ and a positive integer k, the DENSITY SELECTION PROBLEM is to find the consecutive subsequence $A(i^*, j^*)$ over all $O(n^2)$ consecutive subsequences $A(i, j)$ satisfying width constraint $\ell \leq w(i, j) = \sum_{t=i}^{j} w_t \leq u$ such that the rank of its density $d(i^*, j^*) = \sum_{t=i^*}^{j^*} a_t / w(i^*, j^*)$ is k.

The maximum-density segment problem is a special case of the density selection problem such that k is equal to the total number of consecutive subsequences satisfying width constraint and can be solved in optimal $O(n)$ time. The slope selection problem is also a special case of the density selection problem such that $\ell = 0, u = \infty$ and can be solved in optimal $O(n \log n)$ time. The well-known selection problem is also a special case of the density selection problem such that $\ell = 1$, $u = 1$ and $w_i = 1$ for each i and can be solved in optimal $O(n)$ time.

We will give a randomized algorithm for density selection problem that runs in optimal expected $O(n \log n)$ time. On the other hand, we also consider the Density Range Query Problem.

Given a sequence A of n ordered pairs (a_i, w_i) and two width bounds ℓ, u and two real numbers d_l, d_r, the reporting mode of this problem is to find all consecutive subsequences $A(i,j)$ satisfying width constraint such that $d_l \leq d(i,j) \leq d_r$ and the counting mode is just to output the total number of consecutive subsequences $A(i,j)$ satisfying width constraint such that $d_l \leq d(i,j) \leq d_r$. We will show that the reporting mode can be solved in optimal $O(n \log m + h)$ time, where $m = min\{\frac{u-\ell}{w_{min}}, n\}$ and h is the output size, and the counting mode can be solved in optimal $O(n \log m)$ time.

Keywords: Density selection problem, maximum-density segment problem, slope selection problem, selection problem, density range query problem, priority search tree, order-statistics tree, GC content

Joint work with: Lee, D. T.

See also: Proceedings of 20th International Symposium on Algorithms and Computation, Lecture Notes on Computer Science, vol. 5878, 1004-1013, 2009

On the Expected Maximum Degree of Gabriel and Yao Graphs

Pat Morin (Carleton University - Ottawa, CA)

Motivated by applications of Gabriel graphs and Yao graphs in wireless ad-hoc networks, we show that the maximum degree of a random Gabriel graph or Yao graph defined on n points drawn uniformly at random from a unit square grows as $\Theta(\log n / \log \log n)$ in probability.

Keywords: Spanners, probabilistic analysis, Gabriel graph, Yao graph

Joint work with: Devroye, Luc; Gudmundsson, Joachim

Full Paper:
http://arxiv.org/abs/0905.3584

Sparsification Techniques

Yuri Rabinovich (Haifa University - Haifa, IL)

The celebrated Johnson-Lindenstrauss results states that any Euclidean metric on n points can be embedded in $O(\log n/\epsilon^2)$-dimensional space with distortion at most $1+\epsilon$.

The situation for l_1 metrics is quite different from that of l_2 metrics. The results of Brikman-Charikar and Lee-Naor show that one cannot expect dimension reduction better than $d = n^{1/D^2}$ if one is prepared to accept distortion bounded by D. The best upper bound on the dimension, for any constant distortion is $O(n \log n)$, due to Schechtman. The gap between the lower and upper bound is open at large.

We note that by employing results on graph sparsification of Karger-Benczur and Spielmann-Teng-Batson-Srivastava one can show that every l_1 metric can be approximated by embedding into $O(n)$ dimensional l_1 space. We further study these methods and show that they generalize considerably.

In particular, this has leads us to new definitions of finite volume spaces, and their l_1 counterparts. We show that dimension reduction works for such objects too.

Improved Algorithms for Routing and Fully Dynamic Spanners

Liam Roditty (Bar-Ilan University - Ramat Gan, IL)

For a set S of points in a metric space of doubling dimension d, a t-spanner is a sparse graph on the points of S such that between any pair of points there is a path in the spanner whose total length is at most t times the distance between the points. In this paper, we show how to construct a $(1+\epsilon)$-spanner with $n\epsilon^{-O(d)}$ edges and maximum degree $\epsilon^{-O(d)}$ in time $O(n \log n)$. A spanner with similar properties was previously known. However, using our new construction (coupled with several other innovations) we obtain new results for two fundamental problems for constant doubling dimension metrics:

The first result is an essentially optimal compact routing scheme. In particular, we show how to perform routing with a stretch of $1 + \epsilon$, where the label size is $\lceil \log n \rceil$ and the size of the table stored at each point is only $\epsilon^{-O(d)} \log n$. This routing problem was first considered by Peleg and Hassin, who presented a routing scheme in the plane. Later, Chan *et.al.* and Abraham *et.al.* considered this problem for doubling dimension metric spaces. Abraham *et.al.* were the first to present a $(1 + \epsilon)$ routing scheme where the label size depends solely on the number of points. In their scheme labels are of size of $\lceil \log n \rceil$, and each point stores a table of size $\epsilon^{-O(d)} \log^2 n$. In our routing scheme, we achieve routing tables of size $\epsilon^{-O(d)} \log n$, which is essentially the same size as a label (up to the factor of $\epsilon^{-O(d)}$).

The second and main result of this paper is the first **fully** dynamic geometric spanner with poly-logarithmic update time for both insertions and deletions. We present an algorithm that allows points to be inserted into and deleted from S with an amortized update time of $O(\log^3 n)$.

Joint work with: Gottlieb , Lee-Ad

See also: SODA 2008

The Weak Gap Property

Michiel Smid (Carleton University - Ottawa, CA)

We introduce the weak gap property for directed graphs whose vertex set S is a metric space of size n. We prove that, if the doubling dimension of S is a constant, any directed graph satisfying the weak gap property has $O(n)$ edges and total weight $O(\log n) \cdot wt(MST(S))$, where $wt(MST(S))$ denotes the weight of a minimum spanning tree of S. We show that 2-optimal TSP tours and greedy spanners satisfy the weak gap property.

Keywords: Geometric graph, gap property, minimum spanning tree, traveling salesperson problem

Algorithm Engineering for Route Planning in Realistic Scenarios

Dorothea Wagner (KIT - Karlsruhe, DE)

Nowadays, route planning systems belong to the most frequently used information systems. The algorithmic core problem of such systems, i.e., the fast computation of shortest paths is a classical problem that can be solved by Dijkstra's shortest paths algorithm. However, algorithms for route planning in transportation networks have recently undergone a rapid development, leading to methods that are up to three million times faster than Dijkstra's algorithm.

In particular, computing shortest paths in huge networks has become a showpiece of Algorithm Engineering demonstrating the engineering cycle that consists of design, analysis, implementation and experimental evaluation of practicable algorithms.

We will provide a condensed overview of the techniques enabling this development. The main part of the talk will focus on variants of the problem that occur in more realistic traffic scenarios.

Keywords: Algorithm engineering, route planning, shortest paths

Laplace-Beltrami Operator for Spectral Point Clouds Processing

Yusu Wang (Ohio State University, US)

Laplace-Beltrami operator of a manifold is a fundamental mathematical object that encode richs geometric information of the underlying manifold. It has many nice properties. For example, its relation with heat diffusion process makes it a popular tool for smoothing or spectral clustering. Indeed, the Laplace operator (as well as its variants) and its eigenfunctions/eigenvalues have been widely used in a broad range of application areas including machine learning, data mining, and graphics.

In practice, data often comes as a point cloud. In this talk, we consider the case where these points are sampled from a hidden (high-dimensional) manifold. We show how to approximate the Laplacian of the hidden manifold faithfully from such point cloud data (PCD). Previously, theoretical guarantee is only known when the input points are sampled *randomly uniformly* from the hidden manifold, while our result removed this statistical assumption. We remark that this enables us to potentially build a "spectral point clouds processing" framework in non-statistical setting, much like the "spectral graph processing" framework widely used before.

Schematization: A Survey

Alexander Wolff (Universität Würzburg, DE)

In this survey we are interested in (the construction and) the layout of complex networks under angular restrictions concerning network edges. We consider two types of restrictions; edge directions can be restricted to a given discrete set or to a given range. Most of our examples deal with the first type, but the second type also occurs, for example, in underground mining: galleries have a maximum slope. The first type of restriction of edges (to a discrete set of directions) plays a role in orthogonal graph drawing, a recently considered variant called Manhattan-geodesic graph drawing, the layout of Metro maps, path simplification, geographic information systems, (rectilinear and rectangular) cartograms, boundary labeling, VLSI layout, minimum Manhattan networks and other orthogonal spanners.

Keywords: Schematization, orthogonal graph drawing, metro maps, path simplification, GIS, cartograms, VLSI layout, underground mining

Using Bending Invariant Embeddings for Shape Correspondence

Stefanie Wuhrer (NRC - Ottawa, CA)

We discuss a bending invariant representation of a triangular mesh S. The *bending invariant mesh X of S* is an embedding of the intrinsic geometry of S into $\mathbb{R}^3$. That is, the bending invariant mesh X has the property that the geodesic distance between each pair of vertices on S is approximated well by the Euclidean distance between the corresponding vertices on X. Furthermore, X is intersection-free.

Bending invariant meshes can be applied to find dense point-to-point correspondences between a number of isometrically deformed surfaces corresponding to different postures of the same non-rigid object in a fully automatic way.

Joint work with: Shu, Chang; Godin, Guy; Boisvert, Jonathan; Xi, Pengcheng

Stretch Factor of a Plane Geometric Graph in Subquadratic Time

Christian Wulff-Nilsen (University of Copenhagen, DK)

Given a plane geometric graph, its stretch factor is defined as the maximum, over all pairs of vertices, of the ratio between the graph distance and the Euclidean distance between the two vertices. It was an open problem whether the stretch factor can be computed in subquadratic time. We give an algorithm with $O(n^2(\log\log n)^4/\log n)$ running time, where n is the size of the graph, thereby solving the open problem. Our algorithm can easily be adapted to compute the diameter and Wiener index of a planar weighted digraph in subquadratic time, thereby solving two other open problems. The ideas we develop also allow us to obtain an exact distance oracle with subquadratic preprocessing time.

Keywords: Stretch factor, detour, dilation, Wiener index, diameter, planar graph, plane graph

3 Conclusion

We are grateful to the Dagstuhl research center for hosting this successful seminar and for providing such a stimulating and creative environment.